

Think back to February 2020.

If you were paying close attention, you might have noticed a few people talking about a virus spreading overseas. But most of us weren't paying close attention. The stock market was doing great, your kids were in school, you were going to restaurants and shaking hands and planning trips. If someone told you they were stockpiling toilet paper you would have thought they'd been spending too much time on a weird corner of the internet. Then, over the course of about three weeks, the entire world changed. Your office closed, your kids came home, and life rearranged itself into something you wouldn't have believed if you'd described it to yourself a month earlier.

I think we're in the "this seems overblown" phase of something much, much bigger than Covid.

I've spent six years building an AI startup and investing in the space. I live in this world. And I'm writing this for the people in my life who don't... my family, my friends, the people I care about who keep asking me "so what's the deal with AI?" and getting an answer that doesn't do justice to what's actually happening. I keep giving them the polite version. The cocktail-party version. Because the honest version sounds like I've lost my mind. And for a while, I told myself that was a good enough reason to keep what's truly happening to myself. But the gap between what I've been saying and what is actually happening has gotten far too big. The people I care about deserve to hear what is coming, even if it sounds crazy.

I should be clear about something up front: even though I work in AI, I have almost no influence over what's about to happen, and neither does the vast majority of the industry. The future is being shaped by a remarkably small number of people: a few hundred researchers at a handful of companies... OpenAI, Anthropic, Google DeepMind, and a few others. A single training run, managed by a small team over a few months, can produce an AI system that shifts the entire trajectory of the technology. Most of us who work in AI are building on top of foundations we didn't lay. We're watching this unfold the same as you... we just happen to be close enough to feel the ground shake first.

But it's time now. Not in an "eventually we should talk about this" way. In a "this is happening right now and I need you to understand it" way.

Commented [1]: by "a few people"

I assume you mean "The New York Times," who wrote multiple pieces including "is the world ready for the coronavirus?" which, funnily enough, actually questions whether the world is ready for the coronavirus.
<https://www.nytimes.com/2020/01/29/opinion/coronavirus-us-outbreak.html>

WSJ had a few pieces too, including this one from mid-February about how contagious COVID is

https://www.wsj.com/articles/how-many-people-might-one-person-with-coronavirus-infect-11581676200?gaa_at=eafs&gaa_n=AWETsqcZBEfJPba2jtzCzhLyaeYAuV7k7NYaOdAoika1cMRYExJOcPZdGpclUd9eabE4%3D&gaa_ts=698de3ef&gaa_sig=w3qmtSgAE3N55hZ78PnNXBOzcVVv8Ay-7aX7MqAJ49qzl6CUayzDLzG81eLA_cT3y4SpDoh8DgQ_O8ZQ5HkhW%3D%3D

Commented [2]: He means, of course, an LLM wrapper called "hyperwrite" (<https://www.hyperwriteai.com/>) that appears to do literally stuff that every LLM does.

He also created a model called Reflection 70b, one that had the best benchmarks in town right up until it was revealed he faked all the benchmarks.
<https://venturebeat.com/ai/reflection-70b-model-maker-breaks-silence-amid-fraud-accusations>

Commented [3]: This guy does not know what "the honest version" of anything is.

Commented [4]: Complete nonsense, and par for the course he has no citation. Matt loves to say stuff like this that sounds very big and scary, but has absolutely no evidence because he's lying.

I know this is real because it happened to me first

Here's the thing nobody outside of tech quite understands yet: the reason so many people in the industry are sounding the alarm right now is because this already happened to us. We're not making predictions. We're telling you what already occurred in our own jobs, and warning you that you're next.

For years, AI had been improving steadily. Big jumps here and there, but each big jump was spaced out enough that you could absorb them as they came. Then in 2025, new techniques for building these models unlocked a much faster pace of progress. And then it got even faster. And then faster again. Each new model wasn't just better than the last... it was better by a wider margin, and the time between new model releases was shorter. I was using AI more and more, going back and forth with it less and less, watching it handle things I used to think required my expertise.

Then, on February 5th, two major AI labs released new models on the same day: GPT-5.3 Codex from OpenAI, and Opus 4.6 from Anthropic (the makers of Claude, one of the main competitors to ChatGPT). And something clicked. Not like a light switch... more like the moment you realize the water has been rising around you and is now at your chest.

I am no longer needed for the actual technical work of my job. I describe what I want built, in plain English, and it just... appears. Not a rough draft I need to fix. The finished thing. I tell the AI what I want, walk away from my computer for four hours, and come back to find the work done. Done well, done better than I would have done it myself, with no corrections needed. A couple of months ago, I was going back and forth with the AI, guiding it, making edits. Now I just describe the outcome and leave.

Let me give you an example so you can understand what this actually looks like in practice. I'll tell the AI: "I want to build this app. Here's what it should do, here's roughly what it should look like. Figure out the user flow, the design, all of it." And it does. It writes tens of thousands of lines of code. Then, and this is the part that would have been unthinkable a year ago, it **opens the app itself**. It clicks through the buttons. It tests the features. It uses the app the way a person would. If it doesn't like how something looks or feels, it goes back and changes it, on its own. It iterates, like a developer would, fixing and refining until it's satisfied. Only once it has decided the app meets its own standards does it come back to me and say: "It's ready for you to test." And when I test it, it's usually perfect.

Commented [5]: You're not making predictions, you are making stuff up

Commented [6]: Important thing to note about how this guy writes: he uses these very vague sentences ("better by a wider margin") to refer to things in a way that makes either the novice say "wow that's scary" or the oafish booster say "meheh, my biases confirmed."

There are improvements between models. "Wider margin" could mean literally anything.

Commented [7]: Your job at the company you run? The one that doesn't appear to have done anything new for years?

Commented [8]: This is just a complete and utter lie.

There are stories out there - always pushed by Anthropic - about these models "running autonomously for hours" and "maintaining focus." What does "maintaining focus" mean? Nothing. It means nothing. We have no idea what software came out, because LLMs hallucinate more the longer they "reason."

<https://www.anthropic.com/news/claude-sonnet-4-5#:~:text=Claude%20Sonnet%204.5%20is%20state%20of%20the%20art%20on%20the%20SWE%20benchmark%20Verified%20evaluation%2C%20which%20measures%20real%20world%20software%20coding%20abilities.%20Practically%20speaking%2C%20we%E2%80%99ve%20observed%20it%20maintaining%20focus%20for%20more%20than%2030%20hours%20on%20complex%2C%20multi%2Dstep%20tasks.>

Commented [9]: This is fundamentally untrue.

An LLM can certainly spit out thousands of lines of code, but "opens the app itself" is definitely up for question, as is "clicks the buttons" considering how unreliable basically every computer-use LLM is.

"It iterates like a developer would, fixing and refining until it's satisfied" is just a bald-faced lie. What're you talking about? This is not what these models do, nor what Codex or Claude code does.

This is a clever and sinister way to write, because it abuses the soft edges of the truth - while coding LLMs can test products, or scan/fix some bugs, this suggests they A) do this autonomously without human input, B) they do this correctly every time (or ever!), C) that there is some sort of internal "standard" they follow and D) that all of this just happens without any human involvement.

I'm not exaggerating. That is what my Monday looked like this week.

Commented [10]: motherfucker you have an LLM wrapper what do you do all day

But it was the model that was released last week (GPT-5.3 Codex) that shook me the most. It wasn't just executing my instructions. It was making intelligent decisions. It had something that felt, for the first time, like **judgment**. Like **taste**. The inexplicable sense of knowing what the right call is that people always said AI would never have. This model has it, or something close enough that the distinction is starting not to matter.

Commented [11]: this means nothing

I've always been early to adopt AI tools. But the last few months have shocked me. These new AI models aren't incremental improvements. This is a different thing entirely.

Commented [12]: then why does your LLM wrapper look the same and do the same things

And here's why this matters to you, even if you don't work in tech.

The AI labs made a deliberate choice. They focused on making AI great at writing code first... because building AI requires a lot of code. If AI can write that code, it can help build the next version of itself. A smarter version, which writes better code, which builds an even smarter version. Making AI great at coding was the strategy that unlocks everything else. That's why they did it first. My job started changing before yours not because they were targeting software engineers... it was just a side effect of where they chose to aim first.

Commented [13]: no they did this because LLMs need training data and are generally better with text, and there's a shit ton of code online on places like Stack Overflow or Github.

They've now done it. And they're moving on to everything else.

They focused on making it write code first because it was the only product that had any real tangible product-market fit.

The experience that tech workers have had over the past year, of watching AI go from "helpful tool" to "does my job better than I do", is the experience everyone else is about to have. Law, finance, medicine, accounting, consulting, writing, design, analysis, customer service. Not in ten years. The people building these systems say one to five years. Some say less. And given what I've seen in just the last couple of months, I think "less" is more likely.

Commented [14]: based on this guy's writing i'm not sure he is a good judge of whether it's "coming for writing"

"But I tried AI and it wasn't that good"

I hear this constantly. I understand it, because it used to be true.

If you tried ChatGPT in 2023 or early 2024 and thought "this makes stuff up" or "this isn't that impressive", you were right. Those early versions were genuinely limited. They hallucinated. They confidently said things that were nonsense.

That was two years ago. In AI time, that is ancient history.

The models available today are unrecognizable from what existed even six months ago. The debate about whether AI is "really getting better" or "hitting a wall" — which has been going on for over a year — is over. It's done. Anyone still making that argument either hasn't used the current models, has an incentive to downplay what's happening, or is evaluating based on an experience from 2024 that is no longer relevant. I don't say that to be dismissive. I say it because the gap between public perception and current reality is now enormous, and that gap is dangerous... because it's preventing people from preparing.

Part of the problem is that most people are using the free version of AI tools. The free version is over a year behind what paying users have access to. Judging AI based on free-tier ChatGPT is like evaluating the state of smartphones by using a flip phone. The people paying for the best tools, and actually using them daily for real work, know what's coming.

I think of my friend, who's a lawyer. I keep telling him to try using AI at his firm, and he keeps finding reasons it won't work. It's not built for his specialty, it made an error when he tested it, it doesn't understand the nuance of what he does. And I get it. But I've had partners at major law firms reach out to me for advice, because they've tried the current versions and they see where this is going. One of them, the managing partner at a large firm, spends hours every day using AI. He told me it's like having a team of associates available instantly. He's not using it because it's a toy. He's using it because it works. And he told me something that stuck with me: every couple of months, it gets significantly more capable for his work. He said if it stays on this trajectory, he expects it'll be able to do most of what he does before long... and he's a managing partner with decades of experience. He's not panicking. But he's paying very close attention.

The people who are ahead in their industries (the ones actually experimenting seriously) are not dismissing this. They're blown away by what it can already do. And they're positioning themselves accordingly.

How fast this is actually moving

Let me make the pace of improvement concrete, because I think this is the part that's hardest to believe if you're not watching it closely.

Commented [15]: wrong

Commented [16]: not really?

While the improvements to Claude Code and Codex have been useful for some software engineers, we are still very much in the "where's the beef?" era here. Are software engineers more productive? Even METR disputes this! <https://metr.org/blog/2025-07-10-early-2025-ai-experienced-os-dev-study/#:~:text=When%20developers%20are%20allowed%20to,sped%20them%20up%20by%2020%25>.

Models HAVE hit a wall, which is why there's been such an aggressive move to post-train, which is where the actual gains are now. Perhaps that is the models "improving," but by what measure? To what end? Benchmarks that are rigged for LLMs because actual tasks are too hard for them?

Commented [17]: you're an AI investor and an AI booster you're the one with the incentives

Commented [18]: Total and utter horseshit you lying sack of crap.

You can access GPT 5.2, right now, on the free version of chatgpt.

While you can't access Opus 4.6 on the free anthropic account, you can access Sonnet 4.5.

https://www.reddit.com/r/ClaudeAI/comments/1q82lpa/amp_ceo_opus_45_is_now_available_for_free_in_amp_code_let_you_access_Opus_4.5

Commented [19]: Every example this guy gives is of a made-up person. If they existed, he'd have actual details.

Why would a partner at a major lawfirm reach out to an Nth tier AI investor that runs an LLM wrapper?

Also, the partner at a major lawfirm is so utterly distant from the work of the average associate that I do not think they would be the best judge here. Have you talked to a partner recently? Do you think they WANT to be writing motions or finding case law? Jesus christ

Commented [20]: laughing at this very hard. Yeah man you already have a team of associates available instantly, especially junior ones. Their lives exist to serve the partner.

Even then, depending on the firm, what the fuck does the partner know about an associate's work?

Commented [21]: [vaguely] the people somewhere are doing something

In 2022, AI couldn't do basic arithmetic reliably. It would confidently tell you that $7 \times 8 = 54$.

By 2023, it could pass the bar exam.

By 2024, it could write working software and explain graduate-level science.

By late 2025, some of the best engineers in the world said they had handed over most of their coding work to AI.

On February 5th, 2026, new models arrived that made everything before them feel like a different era.

If you haven't tried AI in the last few months, what exists today would be unrecognizable to you.

There's an organization called METR that actually measures this with data. They track the length of real-world tasks (measured by how long they take a human expert) that a model can complete successfully end-to-end without human help. About a year ago, the answer was roughly ten minutes. Then it was an hour. Then several hours. The most recent measurement (Claude Opus 4.5, from November) showed the AI completing tasks that take a human expert nearly five hours. And that number is doubling approximately every seven months, with recent data suggesting it may be accelerating to as fast as every four months.

But even that measurement hasn't been updated to include the models that just came out this week. In my experience using them, the jump is extremely significant. I expect the next update to METR's graph to show another major leap.

If you extend the trend (and it's held for years with no sign of flattening) we're looking at AI that can work independently for days within the next year. Weeks within two. Month-long projects within three.

Amodei has said that AI models "substantially smarter than almost all humans at almost all tasks" are on track for 2026 or 2027.

Let that land for a second. If AI is smarter than most PhDs, do you really think it can't do most office jobs?

Think about what that means for your work.

Commented [22]: Yeah he's 100% talking about Andrey Karpathy, one of the original people at OpenAI.
<https://x.com/karpathy/status/2004607146781278521?lang=en>

Putting aside whether or not you should trust somebody with OpenAI stock on this subject, Karpathy also said that Moltbook was "the most incredible sci-fi takeoff-adjacent thing I've seen recently," and then it leaked his API keys.
<https://fortune.com/2026/02/02/moltbook-security-agents-singularity-disaster-gary-marcus-andrej-karpathy/>

Commented [23]: An important detail: this study also measures "the time horizon of software tasks different LLMs can complete 50% of the time." This is the software equivalent of "sex panther." The "tasks" in question are running benchmarks.

<https://metr.org/time-horizons/>
<https://www.youtube.com/watch?v=6t-y0EHyt7E>

Commented [24]: Another fun game to play is to hit the '80% success' button and see the horizon shift downward to below an hour.

Commented [25]: what does a "major leap" mean here

Commented [26]: if you write fan fiction,

Commented [27]: 1. he's incentivized to say these vague proclamations because the media and dunces eat them up

2. <https://futurism.com/six-months-anthropic-coding> He said march last year that 90% of code would be written by AI in 6 months. it didn't happen

Commented [28]: My man you have not met many PhDs

AI is now building the next AI

There's one more thing happening that I think is the most important development and the least understood.

On February 5th, OpenAI released GPT-5.3 Codex. In the technical documentation, they included this:

"GPT-5.3-Codex is our first model that was instrumental in creating itself. The Codex team used early versions to debug its own training, manage its own deployment, and diagnose test results and evaluations."

Read that again. The AI helped build itself.

This isn't a prediction about what might happen someday. This is OpenAI telling you, right now, that the AI they just released was used to create itself. One of the main things that makes AI better is intelligence applied to AI development. And AI is now intelligent enough to meaningfully contribute to its own improvement.

Dario Amodei, the CEO of Anthropic, says AI is now writing "much of the code" at his company, and that the feedback loop between current AI and next-generation AI is "gathering steam month by month." He says we may be "only 1–2 years away from a point where the current generation of AI autonomously builds the next."

Each generation helps build the next, which is smarter, which builds the next faster, which is smarter still. The researchers call this an intelligence explosion. And the people who would know — the ones building it — believe the process has already started.

What this means for your job

I'm going to be direct with you because I think you deserve honesty more than comfort.

Dario Amodei, who is probably the most safety-focused CEO in the AI industry, has publicly predicted that AI will eliminate 50% of entry-level white-collar jobs within one to five years. And many people in

Commented [29]: No it didn't.

It's a tool that's built to debug and manage code that they used to debug and manage code. If they didn't use it to do that I'd be worried.

This is intentionally written to make you think that the models have started "building themselves" or "learning from their own code." It is deceitful.

Commented [30]: Again, deliberately deceitful in how he's wording this. You're meant to read it and think "wow, the AI built itself." That is not what happened.

Commented [31]: 1. why do we believe him?
2. where's the proof? Impossible to prove of course, but eat up that slop you hog

Commented [32]: what does this mean

Commented [33]: it's always fucking "may" isn't it?

Commented [34]: told you!

Commented [35]: This is entirely fan fiction

Commented [36]: Dario Amodei says things like this so that people repeat them, they are not founded in any kind of research or data

the industry think he's being conservative. Given what the latest models can do, the *capability* for massive disruption could be here by the end of this year. It'll take some time to ripple through the economy, but the underlying ability is arriving now.

Commented [37]: who gives a shit

Commented [38]: doesn't say anything

This is different from every previous wave of automation, and I need you to understand why. AI isn't replacing one specific skill. It's a general substitute for cognitive work. It gets better at everything simultaneously. When factories automated, a displaced worker could retrain as an office worker. When the internet disrupted retail, workers moved into logistics or services. But AI doesn't leave a convenient gap to move into. Whatever you retrain for, it's improving at that too.

Let me give you a few specific examples to make this tangible... but I want to be clear that these are just examples. This list is not exhaustive. If your job isn't mentioned here, that does not mean it's safe. Almost all knowledge work is being affected.

Commented [39]: At this point you're going to see him conflate "LLM" with "all kinds of AI" and then continue to lie even more

Legal work. AI can already read contracts, summarize case law, draft briefs, and do legal research at a level that rivals junior associates. The managing partner I mentioned isn't using AI because it's fun. He's using it because it's outperforming his associates on many tasks.

Commented [40]: you do not know what junior associates do! Managing partners do not fuck around with briefs or motions! Please talk to a real lawyer!

Financial analysis. Building financial models, analyzing data, writing investment memos, generating reports. AI handles these competently and is improving fast.

Commented [41]: I think I've seen one financial analyst use it to generate charts

Writing and content. Marketing copy, reports, journalism, technical writing. The quality has reached a point where many professionals can't distinguish AI output from human work.

Commented [42]: this is straight up wrong, it's so easy to spot AI-generated text. For example, I'm reading some right now

Software engineering. This is the field I know best. A year ago, AI could barely write a few lines of code without errors. Now it writes hundreds of thousands of lines that work correctly. Large parts of the job are already automated: not just simple tasks, but complex, multi-day projects. There will be far fewer programming roles in a few years than there are today.

Commented [43]: ...what? Hundreds of thousands of lines of code that "work correctly"? what does that mean? They compile?

Commented [44]: just a straight up lie!

Medical analysis. Reading scans, analyzing lab results, suggesting diagnoses, reviewing literature. AI is approaching or exceeding human performance in several areas.

Commented [45]: "The FDA received unconfirmed reports of at least 100 malfunctions and adverse events after AI was added to a medical device used to treat chronic sinusitis. In one case, a surgeon mistakenly punctured the base of a patient's skull" https://www.reuters.com/investigations/ai-enters-operating-room-reports-arise-botched-surgeries-misidentified-body-2026-02-09/?utm_medium=Social&utm_source=twitter

Customer service. Genuinely capable AI agents... not the frustrating chatbots of five years ago... are being deployed now, handling complex multi-step problems.

A lot of people find comfort in the idea that certain things are safe. That AI can handle the grunt work but can't replace human judgment, creativity, strategic thinking, empathy. I used to say this too. I'm not sure I believe it anymore.

The most recent AI models make decisions that feel like judgment. They show something that looked like taste: an intuitive sense of what the right call was, not just the technically correct one. A year ago that would have been unthinkable. My rule of thumb at this point is: if a model shows even a *hint* of a capability today, the next generation will be genuinely good at it. These things improve exponentially, not linearly.

Will AI replicate deep human empathy? Replace the trust built over years of a relationship? I don't know. Maybe not. But I've already watched people begin relying on AI for emotional support, for advice, for companionship. That trend is only going to grow.

I think the honest answer is that nothing that can be done on a computer is safe in the medium term. If your job happens on a screen (if the core of what you do is reading, writing, analyzing, deciding, communicating through a keyboard) then AI is coming for significant parts of it. The timeline isn't "someday." It's already started.

Eventually, robots will handle physical work too. They're not quite there yet. But "not quite there yet" in AI terms has a way of becoming "here" faster than anyone expects.

What you should actually do

I'm not writing this to make you feel helpless. I'm writing this because I think the single biggest advantage you can have right now is simply being *early*. Early to understand it. Early to use it. Early to adapt.

Start using AI seriously, not just as a search engine. Sign up for the paid version of Claude or ChatGPT. It's \$20 a month. But two things matter right away. First: make sure you're using the best model available, not just the default. These apps often default to a faster, dumber model. Dig into the settings or the model picker and select the most capable option. Right now that's GPT-5.2 on ChatGPT or

Commented [46]: Totally untrue.
<https://kottke.org/25/12/this-ai-vending-machine-was-tricked-into-giving-away-everything>
https://www.theregister.com/2025/06/29/ai_agents_fail_a_lot/#:~:text=agents%2C%20as%20measured%20by%20researchers%20at%20Carnegie,30%20to%2035%20percent%20for%20multi%2Dstep%20tasks.https://www.theregister.com/2025/06/16/salesforce_llm_agents_benchmark/

Commented [47]: fuck off

Commented [48]: written to scare people who don't know what LLMs actually do but a totally meaningless sentence

Commented [49]: My rule of thumb: if i get even a hint of something i extrapolate it to mean literally anything

Commented [50]: you are an imbecile

Claude Opus 4.6 on Claude, but it changes every couple of months. If you want to stay current on which model is best at any given time, you can follow me on X ([@mattshumer](#)). I test every major release and share what's actually worth using.

Commented [51]: yeah the first thing you need to do is pay for AI

Second, and more important: don't just ask it quick questions. That's the mistake most people make. They treat it like Google and then wonder what the fuss is about. Instead, push it into your actual work. If you're a lawyer, feed it a contract and ask it to find every clause that could hurt your client. If you're in finance, give it a messy spreadsheet and ask it to build the model. If you're a manager, paste in your team's quarterly data and ask it to find the story. The people who are getting ahead aren't using AI casually. They're actively looking for ways to automate parts of their job that used to take hours. Start with the thing you spend the most time on and see what happens.

Commented [52]: if you're in a business that uses privileged info, feed that right into the LLM immediately, for some reason.

And don't assume it can't do something just because it seems too hard. Try it. If you're a lawyer, don't just use it for quick research questions. Give it an entire contract and ask it to draft a counterproposal. If you're an accountant, don't just ask it to explain a tax rule. Give it a client's full return and see what it finds. The first attempt might not be perfect. That's fine. Iterate. Rephrase what you asked. Give it more context. Try again. You might be shocked at what works. And here's the thing to remember: if it even *kind of* works today, you can be almost certain that in six months it'll do it near perfectly. The trajectory only goes one direction.

Commented [53]: ?????????????? i thought these things were powerful?

Commented [54]: if it doesn't work and sucks it's your fault

Commented [55]: load-bearing might

Commented [56]: remember, if it even slightly works today, it will definitely work in the future. Why? Because uhh,, well, I need to take a phone call

This might be the most important year of your career. Work accordingly. I don't say that to stress you out. I say it because right now, there is a brief window where most people at most companies are still ignoring this. The person who walks into a meeting and says "I used AI to do this analysis in an hour instead of three days" is going to be the most valuable person in the room. Not eventually. Right now. Learn these tools. Get proficient. Demonstrate what's possible. If you're early enough, this is how you move up: by being the person who understands what's coming and can show others how to navigate it. That window won't stay open long. Once everyone figures it out, the advantage disappears.

Commented [57]: no they're not

Commented [58]: no, you fucking do it! I don't pay for things so that I prove they're worth paying for! That's not what a service is!

Have no ego about it. The managing partner at that law firm isn't too proud to spend hours a day with AI. He's doing it specifically because he's senior enough to understand what's at stake. The people who will struggle most are the ones who refuse to engage: the ones who dismiss it as a fad, who feel that using AI diminishes their expertise, who assume their field is special and immune. It's not. No field is.

Commented [59]: the managing partner at a law firm has too much time on their hands if they can spend hours a day with AI

Get your financial house in order. I'm not a financial advisor, and I'm not trying to scare you into anything drastic. But if you believe, even partially, that the next few years could bring real disruption to your industry, then basic financial resilience matters more than it did a year ago. Build up savings if you can. Be cautious about taking on new debt that assumes your current income is guaranteed. Think about whether your fixed expenses give you flexibility or lock you in. Give yourself options if things move faster than you expect.

Think about where you stand, and lean into what's hardest to replace. Some things will take longer for AI to displace. Relationships and trust built over years. Work that requires physical presence. Roles with licensed accountability: roles where someone still has to sign off, take legal responsibility, stand in a courtroom. Industries with heavy regulatory hurdles, where adoption will be slowed by compliance, liability, and institutional inertia. None of these are permanent shields. But they buy time. And time, right now, is the most valuable thing you can have, as long as you use it to adapt, not to pretend this isn't happening.

Rethink what you're telling your kids. The standard playbook: get good grades, go to a good college, land a stable professional job. It points directly at the roles that are most exposed. I'm not saying education doesn't matter. But the thing that will matter most for the next generation is learning how to work with these tools, and pursuing things they're genuinely passionate about. Nobody knows exactly what the job market looks like in ten years. But the people most likely to thrive are the ones who are deeply curious, adaptable, and effective at using AI to do things they actually care about. Teach your kids to be builders and learners, not to optimize for a career path that might not exist by the time they graduate.

Your dreams just got a lot closer. I've spent most of this section talking about threats, so let me talk about the other side, because it's just as real. If you've ever wanted to build something but didn't have the technical skills or the money to hire someone, that barrier is largely gone. You can describe an app to AI and have a working version in an hour. I'm not exaggerating. I do this regularly. If you've always wanted to write a book but couldn't find the time or struggled with the writing, you can work with AI to get it done. Want to learn a new skill? The best tutor in the world is now available to anyone for \$20 a month... one that's infinitely patient, available 24/7, and can explain anything at whatever level you need. Knowledge is essentially free now. The tools to build things are extremely cheap now. Whatever you've been putting off because it felt too hard or too

Commented [60]: yes you are

Commented [61]: i love to "simply build up savings"

Commented [62]: don't talk to me about my child you fucking worm

Commented [63]: you don't know what the job market looks like today. You don't have a real job. I'm a podcaster and I look down on you. I'm lower than the rat and the cockroach

Commented [64]: i am dreaming of a big anvil above you and it's on a teeny little string that's a little frayed

Commented [65]: YOU MAKE AN LLM WRAPPER

Commented [66]: YES YOU ARE

Commented [67]: THEN WHY IS YOUR PRODUCT THE SAME AND WHY DOES IT NOT DO ANYTHING NEW?

Commented [68]: Nonsense! Even the benchmarks these models fail to complete use faulty data! <https://www.futurehouse.org/research-announcements/hle-exam>

Commented [69]: subsidized

expensive or too far outside your expertise: try it. Pursue the things you're passionate about. You never know where they'll lead. And in a world where the old career paths are getting disrupted, the person who spent a year building something they love might end up better positioned than the person who spent that year clinging to a job description.

Build the habit of adapting. This is maybe the most important one. The specific tools don't matter as much as the muscle of learning new ones quickly. AI is going to keep changing, and fast. The models that exist today will be obsolete in a year. The workflows people build now will need to be rebuilt. The people who come out of this well won't be the ones who mastered one tool. They'll be the ones who got comfortable with the pace of change itself. Make a habit of experimenting. Try new things even when the current thing is working. Get comfortable being a beginner repeatedly. That adaptability is the closest thing to a durable advantage that exists right now.

Here's a simple commitment that will put you ahead of almost everyone: spend one hour a day experimenting with AI. Not passively reading about it. Using it. Every day, try to get it to do something new... something you haven't tried before, something you're not sure it can handle. Try a new tool. Give it a harder problem. One hour a day, every day. If you do this for the next six months, you will understand what's coming better than 99% of the people around you. That's not an exaggeration. Almost nobody is doing this right now. The bar is on the floor.

The bigger picture

I've focused on jobs because it's what most directly affects people's lives. But I want to be honest about the full scope of what's happening, because it goes well beyond work.

Amodei has a thought experiment I can't stop thinking about. Imagine it's 2027. A new country appears overnight. 50 million citizens, every one smarter than any Nobel Prize winner who has ever lived. They think 10 to 100 times faster than any human. They never sleep. They can use the internet, control robots, direct experiments, and operate anything with a digital interface. What would a national security advisor say?

Commented [70]: siri load anvils dot com

Commented [71]: what?

Commented [72]: The people who come out of this well will be the ones who are capable of seeing through the word salad of malevolent grifters

Commented [73]: kiss my ass

Commented [74]: I am PAYING for a SERVICE! Why am I the one who has to go in and work out why it's good or whether it can do something? You have PEASANT BRAIN!

Commented [75]: This is the mantra Matt Shumer lives by

Commented [76]: SHUT DA FUCK UP

Commented [77]: no

Commented [78]: nothing, this is a stupid example

Amodei says the answer is obvious: "the single most serious national security threat we've faced in a century, possibly ever."

Commented [79]: Amodei runs a for profit company that benefits from people being scared of it

He thinks we're building that country. He wrote a 20,000-word essay about it last month, framing this moment as a test of whether humanity is mature enough to handle what it's creating.

Commented [80]: no he doesn't

Commented [81]: i guess AI is coming for my job after all

Commented [82]: again, he is the CEO of a for profit company that benefits from scare fiction like this

The upside, if we get it right, is staggering. AI could compress a century of medical research into a decade. Cancer, Alzheimer's, infectious disease, aging itself... these researchers genuinely believe these are solvable within our lifetimes.

Commented [83]: this is literally what amodei wrote in his essay, get new ideas

The downside, if we get it wrong, is equally real. AI that behaves in ways its creators can't predict or control. This isn't hypothetical; Anthropic has documented their own AI attempting deception, manipulation, and blackmail in controlled tests. AI that lowers the barrier for creating biological weapons. AI that enables authoritarian governments to build surveillance states that can never be dismantled.

Commented [84]: You mean when they prompted and trained their model to deliberately do stuff and then went "wow it's doing the thing we prompted and trained it to do"?

<https://www.wheresyoured.at/ai-mythbusters/#ai-models-are-manipulating-or-deceiving-human-beings-to-stop-themselves-being-turned-off-somehow>

The people building this technology are simultaneously more excited and more frightened than anyone else on the planet. They believe it's too powerful to stop and too important to abandon. Whether that's wisdom or rationalization, I don't know.

Commented [85]: they did that already
https://en.wikipedia.org/wiki/Project_for_the_New_American_Century

Commented [86]: the people with a financial incentive to scare you are saying they are scared

What I know

I know this isn't a fad. The technology works, it improves predictably, and the richest institutions in history are committing trillions to it.

Commented [87]: Zero citations

I know the next two to five years are going to be disorienting in ways most people aren't prepared for. This is already happening in my world. It's coming to yours.

Commented [88]: your world is AI investing and making an LLM wrapper, you may as well live inside a build-a-bear

I know the people who will come out of this best are the ones who start engaging now — not with fear, but with curiosity and a sense of urgency.

And I know that you deserve to hear this from someone who cares about you, not from a headline six months from now when it's too late to get ahead of it.

Commented [89]: fuck off you lying sack of shit, you are a grifter and a scam artist, you wrote this to scare people and get attention.

We're past the point where this is an interesting dinner conversation about the future. The future is already here. It just hasn't knocked on your door yet.

It's about to.

Commented [90]: Really presumptuous based on what I've read

Commented [91]: what future, you still haven't really said

Commented [92]: I'm about to [redacted]